Sparse geometries handling in lattice-Boltzmann method implementation for graphic processors

Tadeusz Tomczak, Roman G. Szafran

arXiv:1703.08015v1 [cs.DC] 23 Mar 2017

Abstract—We describe a high-performance implementation of the lattice-Boltzmann method (LBM) for sparse geometries on graphic processors. In our implementation we cover the whole geometry with a uniform mesh of small tiles and carry out calculations for each tile independently with a proper data synchronization at tile edges. For this method we provide both the theoretical analysis of complexity and the results for real implementations for 2D and 3D geometries. Based on the theoretical model, we show that tiles offer significantly smaller bandwidth overhead than solutions based on indirect addressing. For 2-dimensional lattice arrangements a reduction of memory usage is also possible, though at the cost of diminished performance. We reached the performance of 682 MLUPS on GTX Titan (72% of peak theoretical memory bandwidth) for D3Q19 lattice arrangement and double precision data.

Index Terms—GPU, CUDA, LBM, CFD, parallel computing

1 Introduction

The lattice-Boltzmann method (LBM) is a highly parallel alternative to classical Navier-Stokes solvers for computational fluid dynamics (CFD). Many researchers have shown its excellent absolute performance and scalability on modern parallel machines, especially graphic processing units (GPU). A thorough review of GPU LBM implementations is presented in [1].

In many areas of CFD application (e.g. biomedical or porous media simulations) the geometry is sparse (contains many areas without fluid). Although many optimization techniques for LBM implementations on GPU are known [1]–[11], they are not specialized for sparse geometries. Simulations for sparse geometries can be run using methods designed for large-scale multi-GPU implementations [12]–[14], but the lack of specific optimizations can significantly increase memory usage and computational effort required for simulation, what decreases the domain size that can be simulated on a single machine, and increases the number of processors required for a domain of a given size. This can be especially unfavorable for GPUs, where the memory amount per processor is fixed and usually smaller than in typical CPU configurations. Thus, efficient handling of sparse geometries is crucial to achieving a high performance and a good hardware utilization in many real-world simulations.

Current GPU implementations of LBM specialized for sparse geometries presented in [15]–[17] are based on indirect addressing. This limits memory usage but causes significantly larger memory traffic than in dense implementations. For example, a highly optimized sparse implementation from [16] achieves 64% of performance shown in [10] for the same collision model and hardware (Tesla K20).

In this work we present the GPU implementation of LBM for sparse geometries where the information about geometry is stored using the uniform grid of small, fixed-size tiles. The use of tiles reduces to almost negligible values a number of ancillary data transfers needed to manage the geometry sparsity. For this method we provide the theoretical analysis, which allows to find the performance limits and compare them with indirect addressing solutions. We also show the results for two real implementations: a two-dimensional with minimized memory usage and a three-dimensional with high performance.

The structure of the paper is as follows. In section 2 we briefly introduce the lattice-Boltzmann method and analyze existing techniques of sparse geometries handling. Section 3 contains the description of our implementation with the detailed analysis of introduced overheads. The performance comparison with existing implementations is presented in section 4. Section 5 contains conclusions.

2 Lattice-Boltzmann method

2.1 Basics

In the LBM, as in the conventional CFD, the geometry, initial and boundary conditions must be specified to solve the initial-value problem. The computational domain is uniformly partitioned and computational nodes are placed in vertexes of adjacent cells, giving the lattice. One of the lattice structures used in this study, D2Q9, is presented in Fig. 1. The first number in the notation (D2) is the space dimension; the second one (Q9) is the lattice links number.

Let f_i represent the probability distribution function (PDF) along the i lattice direction, δ_x and δ_t are the lattice spacing and the lattice time step respectively, $\frac{\delta_x}{\delta_t}$ is the lattice speed, $\boldsymbol{e}_i$ is the unit vector along the i lattice direction,

- *T. Tomczak is with the Wroclaw University of Science and Technology, Faculty of Electronics, Chair of Computer Engineering, ul. Janiszewskiego 11/17, 50-370 Wroclaw, Poland. E-mail: tadeusz.tomczak@pwr.edu.pl*
- *Roman G. Szafran is with the Wroclaw University of Science and Technology, Faculty of Chemistry, Department of Chemical Engineering, ul. Norwida 4/6, 50-373 Wroclaw, Poland. E-mail:roman.szafran@pwr.edu.pl, phone: +48713203813.*

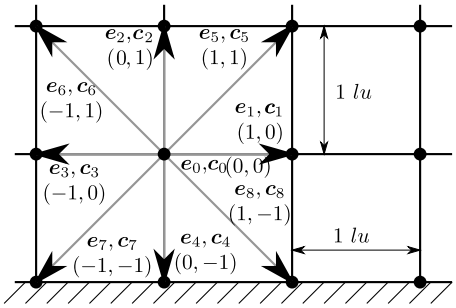

Fig. 1. The D2Q9 lattice structure; $\boldsymbol{e}_i$, $\boldsymbol{c}_i$ are vectors along the i lattice direction; lu denotes the lattice unit.

and $\boldsymbol{c}_i = \frac{\delta_x}{\delta_t}\boldsymbol{e}_i$ is the lattice velocity vector in the velocity space along the i lattice direction. The core of the lattice-Boltzmann method is the discretized in velocity $\boldsymbol{R}^n$ space form of the Boltzmann transport equation that can be written for each i of q directions (lattice linkages) in velocity space as follows:

$$\frac{\partial f_i}{\partial t} + \boldsymbol{c}_i \cdot \nabla f_i = \Omega_i, \tag{1}$$

where Ω_i is the collision operator.

The collision operator can be approximated with the most popular Bhatnagar-Gross-Krook model,

$$\Omega_{BGK} = -\frac{1}{\tau}\left(f_i - f_i^{eq}\right), \tag{2}$$

where τ is the relaxation time in LB units related to the lattice fluid viscosity: $v = 1/3(\tau - 1/2)$ and f_i^{eq} is the equilibrium distribution function along the i lattice direction given by the following formula for quasi-compressible model:

$$f_i^{eq} = \omega_i \rho \left(1 + \frac{\boldsymbol{c}_i \cdot \boldsymbol{u}}{c_s^2} + \frac{(\boldsymbol{c}_i \cdot \boldsymbol{u})^2}{2c_s^4} - \frac{\boldsymbol{u}^2}{2c_s^2}\right) \tag{3}$$

and for incompressible model:

$$f_i^{eq} = \omega_i \left(\rho + \frac{\boldsymbol{c}_i \cdot \boldsymbol{u}}{c_s^2} + \frac{(\boldsymbol{c}_i \cdot \boldsymbol{u})^2}{2c_s^4} - \frac{\boldsymbol{u}^2}{2c_s^2}\right) \tag{4}$$

where $c_s = 1/\sqrt{3}$ is the lattice speed of sound that is a lattice constant; $\boldsymbol{u}$ is the macroscopic fluid velocity vector expressed in LB units; $\rho = \sum_i f_i$ is a fluid density that is related to a pressure $p = \rho/3$, both expressed in LB units, and ω_i is a weighting scalar for the i lattice direction. The macroscopic velocity for quasi-compressible model can be determined from

$$\boldsymbol{u} = \frac{1}{\rho}\sum_i \boldsymbol{c}_i f_i \tag{5}$$

and for incompressible model from

$$\boldsymbol{u} = \sum_i \boldsymbol{c}_i f_i. \tag{6}$$

Integrating Eqn. (1) from t to $t + \delta_t$ along the i lattice direction and assuming that the collision term is constant during the interval, we can obtain discretized in time form of BGK-LBM equation

$$\underbrace{f_i\left(\boldsymbol{r} + \boldsymbol{c}_i\delta_t, t + \delta_t\right) - f_i(\boldsymbol{r}, t)}_{Streaming} = \underbrace{\frac{\delta_t}{\tau}\left[f_i^{eq}(\boldsymbol{r}, t) - f_i(\boldsymbol{r}, t)\right]}_{Collision}, \tag{7}$$

where $\boldsymbol{r}$ is a position vector in the velocity space. The therm on the left hand side is known as the streaming (propagation) step; the latter represents the collision step. These two steps are repeated sequentially during the simulation giving velocity, density and pressure distributions at each time step.

The BGK-LBM uses the single relaxation time to characterize the collision effects. However, physically, these rates should be different during collision processes. To overcome this limitation, a collision matrix with different eigenvalues or multiple relaxation times can be used. The LBM with a MRT collision operator can be expressed as:

$$\underbrace{f_i\left(\boldsymbol{r} + \boldsymbol{c}_i\delta_t, t + \delta_t\right) - f_i(\boldsymbol{r}, t)}_{Streaming} = \underbrace{\boldsymbol{A}\left[\boldsymbol{f}^{eq} - \boldsymbol{f}\right]}_{Collision}, \tag{8}$$

where $\boldsymbol{A}$ is the collision matrix. In this study we use both of the above mentioned collision models in compressible and quasi-compressible versions.

2.2 Performance model

To compare different implementations, we employ a widely used, simple computational complexity model based on [18]. In this model, we assume that the data from device memory is read only once and cached in fast, internal memory (registers/cache/scratchpad), and that the device memory can be effectively read in single-byte transactions. These assumptions are rather unrealistic for current machines because both CPU and GPU use DRAM memories with burst transactions and have the limited size of registers and cache. However, such "ideal" machine model allows finding the minimal amounts of data that have to be stored and transferred during computations. The minimal amounts of data define both the minimum memory usage and the maximum achievable performance for bandwidth bound implementations.

For simplification, we use the same measure of complexity for fluid and boundary nodes. In many geometries the boundary nodes are only a small portion of all non-solid nodes. Moreover, many boundary nodes do not require significantly different amounts of stored and transferred data. Even if additional values are used, they are often shared between many nodes and can be placed in cache/constant memory.

Let q denote the number of PDF f_i and s_d denote the number of bytes for storing a single f_i value (e.g. 4/8 B for a single/double precision floating point). Thus, the minimum number of memory bytes needed for single node datum is

$$M_{node} = q \cdot s_d \ \ [B] \tag{9}$$

and the minimum number of bytes transferred per one LBM iteration for a single node is

$$B_{node} = 2 \cdot q \cdot s_d \ \ [B]. \tag{10}$$

Notice that we assume that an information about the node type is not stored in memory.

Formulas defining the number of floating point operations (FLOP, not to be confused with floating point operations per second, FLOPS) are much more complex to determine. A simple count of the operations resulting from a naive implementation of equations (7) and (8) gives numbers significantly larger than in real implementations due to optimizations carried out at compilation stage (multiplications by constants $e_i \in \{-1, 0, 1\}$, common subexpression elimination, constant folding etc.). Thus, for numbers of computational operations, we show only specific values obtained by a disassemble of a GPU binary code using the *nvdisasm* utility (we count only floating point arithmetic operations, and fused multiply-add is treated as two operations). The complete number of FLOP per D2Q9 lattice node (including computations of velocity and density) varies from 52 FLOP for BGK incompressible to 145 FLOP for MRT quasi-compressible, where the floating point inverse must be computed on a GPU. For D3Q19 lattice the computations require between 304 FLOP for BGK incompressible and 1165 FLOP for MRT quasi-compressible.

The numbers of FLOP cannot be treated as the minimal numbers of operations but can be used to estimate whether a performance of LBM implementation on a given machine is bound by memory bandwidth, instruction throughput, or latency. According to Eqn. (10) each node requires the transfer of 144 bytes for D2Q9 and 304 bytes for D3Q19 lattice arrangements and double precision data. This gives between 2.77 B/FLOP (BGK incompressible) and 0.99 B/FLOP (MRT quasi-compressible) for D2Q9 and between 1 B/FLOP (BGK incompressible) and 0.26 B/FLOP (MRT quasi-compressible) for D3Q19 lattice arrangement. Comparing these numbers with defined in [19] the *machine balance* for GPU (about 0.29 B/FLOP for Fermi, 0.18 B/FLOP for Kepler, and 0.14 B/FLOP for Pascal architectures) it can be seen that LBM implementations on GPU should be usually bandwidth-bound. Only the performance of MRT quasi-compressible implementation for D3Q19 lattice arrangement on Fermi architectures can be limited by arithmetic operations.

2.3 Sparsity handling overhead

Compared with the minimal requirements defined in section 2.2 the techniques for sparse geometries bring some memory and bandwidth overheads due to indirect addressing. We ignore the computational overhead because the LBM implementations on GPU are usually bandwidth bound. Let the overhead Δ (Δ^M - memory overhead, Δ^B - bandwidth overhead) be defined as a ratio of additional operations to the minimum numbers defined in Eqn. (9) and (10). Only a single LBM time iteration is analyzed since all iterations are processed in the same way.

Let N_{snodes} and N_{fnodes} denote the number of solid and non-solid (fluid and boundary) nodes in a geometry. The number of all nodes in the geometry is then $N_{nodes} = N_{snodes} + N_{fnodes}$. We can also define a geometry porosity

$$\phi = \frac{N_{snodes}}{N_{nodes}} \tag{11}$$

and a solidity

$$\eta = \frac{N_{fnodes}}{N_{nodes}}, \tag{12}$$

where both $\phi, \eta \in [0, 1]$ and $\phi + \eta = 1$.

2.3.1 Connectivity Matrix

The GPU implementations of Connectivity Matrix (CM), shown in [15] and an optimized version in [16], use pointers to the neighbor node for each propagated function f_i. No data are stored for solid nodes, but the additional pointers for each f_i function are needed for each non-solid node. Since the pointer can be replaced by an index to a linear array, then for GPU implementations the $s_{CM_idx} = 4$ byte index is enough (6GB allows for storing less than $200 \cdot 10^6$ nodes for the D2Q9 lattice and a single copy of single precision f_i values). Thus, each non-solid node requires an additional storage of $q - 1$ indexes (q for implementation from [16]) and the memory overhead is

$$\Delta^M_{CM} = \frac{(q-1) \cdot s_{CM_idx}}{M_{node}} + 1, \tag{13}$$

where the "+1" term denotes the second copy of f_i data for each non-solid node used to avoid race conditions in both [15] and [16]. For double precision f_i values and $s_{CM_idx} = 4$ bytes the left term in Eqn. (13) is less than 0.5.

The bandwidth overhead for CM results only from reads of indexes from CM. Thus, using the minimum transfer for a single node defined in Eqn. (10), the bandwidth overhead is

$$\Delta^B_{CM} = \frac{(q-1) \cdot s_{CM_idx}}{B_{node}} \approx \frac{s_{CM_idx}}{2 \cdot s_d}. \tag{14}$$

The value of Δ^B_{CM} depends mainly on the ratio of s_{CM_idx} to s_d, and is < 0.5 for single, and < 0.25 for double precision f_i values.

2.3.2 Fluid Index Array

The GPU implementation of fluid index array (FIA) technique, presented in [17], uses a "bitmap" that for each node contains either a pointer to the non-solid node data or -1 when a node is solid. The values of f_i functions are stored only for non-solid nodes. Thus, the additional memory is used only for pointers from FIA (single pointer per each solid/non-solid node). Similarly as for CM, pointers can be replaced with $s_{FIA_idx} = 4$ B indexes. The memory overhead for FIA is

$$\Delta^M_{FIA} = \frac{s_{FIA_idx} \cdot N_{nodes}}{N_{fnodes} \cdot M_{node}} + 1 = \frac{s_{FIA_idx}}{\eta \cdot M_{node}} + 1, \tag{15}$$

where the "+1" term denotes the second copy of f_i data for each non-solid node used in [17] to avoid race conditions.

For geometries with a low number of solid nodes the left term from Eqn. (15) is much smaller than 1 (about $\frac{0.111}{\eta}$ for single and about $\frac{0.056}{\eta}$ for double precision f_i values and D2Q9 lattice arrangement). However, when a geometry is very sparse, then the value of $\frac{s_{FIA_idx}}{\eta \cdot M_{node}}$ grows fast: for D2Q9 lattice arrangement and single precision f_i values it exceeds 1.0 for $\eta = 0.1$.

The bandwidth overhead for FIA based implementation is more difficult to define. Fast LBM implementations fuse

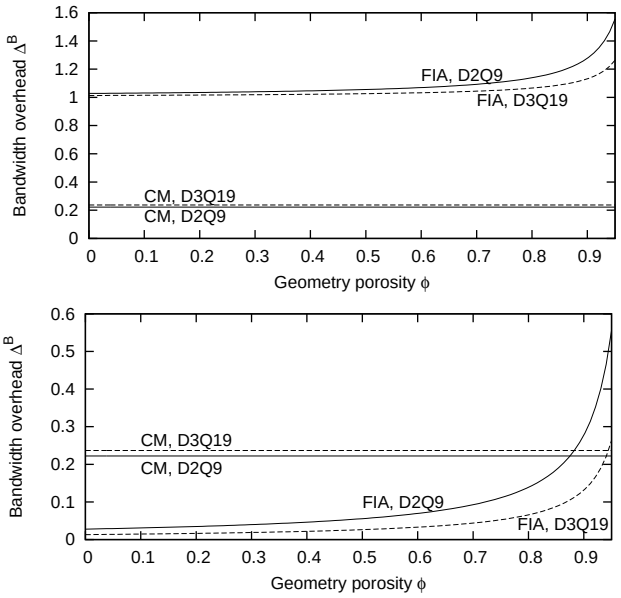

Fig. 2. Comparison of bandwidth overhead estimates for real (top) and ideal (bottom) implementations of CM and FIA techniques. All results for double precision f_i values and $s_{CM_idx} = s_{FIA_idx} = 4$ bytes.

collision and streaming steps into a single kernel. Unfortunately, the original implementation from [17] uses two separate kernels. The first kernel, launched only for fluid nodes, reads and writes f_i values and realizes collision step. The second kernel, launched for all nodes (also the solid ones), is responsible for the streaming step and in addition to reads and writes of f_i values also reads the FIA indexes for the current node and for all neighbor nodes if the current node is non-solid. This approach significantly increases bandwidth overhead (to values larger than 1) due to double access to f_i values. In our estimation, we treat the additional reads/writes of f_i values as required for FIA based implementations and model these by adding the "+1" term to the bandwidth overhead. The rest of the overhead results only from additional reads of indexes from FIA. Since we assume that the indexes from FIA are read from device memory only once and cached internally, the bandwidth overhead is

$$\Delta^B_{FIA} = \frac{N_{nodes} \cdot s_{FIA_idx}}{N_{fnodes} \cdot B_{node}} + 1 = \frac{s_{FIA_idx}}{\eta \cdot B_{node}} + 1. \qquad (16)$$

Notice that $\frac{s_{FIA_idx}}{\eta \cdot B_{node}}$ equals to half of the left term from Eqn. (15).

2.3.3 Ideal implementations

The overhead estimates shown in Eqn. (13)–(16) take into account the characteristics of real implementations, what results in the additional "+1" terms in equations for Δ^M_{CM}, Δ^M_{FIA} and Δ^B_{FIA}. We can also consider the "ideal" implementations (with these "+1" terms removed in some way) that can be used to define the performance limits of CM/FIA techniques.

The comparison of estimated bandwidth overheads for CM and FIA techniques is shown in Fig. 2. Although the real implementation of FIA has 4 times higher bandwidth overhead than CM, in theory the ideal FIA implementation can result in lower bandwidth overhead than CM for geometries with porosity lower than 0.875 (0.944) for D2Q9 (D3Q19) lattice arrangement.

The plots of estimated memory overheads for CM and FIA have a shape similar to shown in Fig. 2 bottom for both real and ideal implementations. Notice that the overheads for FIA decrease with an increase of the lattice links number.

The bandwidth overheads for ideal implementations of FIA and CM strongly depend on s_{FIA_idx} and s_{CM_idx}. By decreasing the size of indexes to 2 or 1 B the overhead can be decreased twice or four times. However, decreasing the size of the neighbor node index seems not trivial for large geometries.

2.3.4 Other methods

Simulations for sparse geometries can be also run using codes for dense geometries, usually on multi-GPU clusters due to large memory requirements. Depending on the applied solutions, the memory and performance penalties caused by geometry sparsity can vary.

For the simplest case all data must be stored and transferred for all nodes (fluid and solid), thus $1/\eta$ times more memory and transfers are needed. For example, storing data for all nodes from a geometry with $\eta = 0.1$ requires ten times more memory than storing data only for fluid nodes. Since FIA/CM require about 2–$2.5\times$ more memory per node than implementations for dense geometries, thus the use of code designed for dense geometries to simulate the sparse ones requires about 4–$5\times$ more memory when $\phi = 0.9$.

The bandwidth overhead of implementation for dense geometries can be greatly reduced through skipping of operations for solid nodes after checking a type of node. This optimization can reduce the bandwidth overhead to values lower even than for ideal implementations of FIA because usually the node type can be encoded on a smaller number of bytes than s_{FIA_idx}. However, this does not take into account the uncoalesced memory transactions resulting from interlacing in memory data for solid and non-solid nodes.

The interesting technique is presented in waLBerla, the large-scale multi-GPU framework [12]–[14], where the geometry is divided using the hierarchical structure of "Patches" composed of "Blocks". For sparse geometries empty blocks can be removed reducing memory usage and computational complexity. Though tiles presented in this work use a similar concept, Blocks are a higher level data structure designed rather for efficient multi-processor implementations, where a load balancing and a communication affect performance. The main differences between tiles and Blocks are size (tiles have small, fixed size, additionally limited by hardware parameters, for example warp/memory transaction/cache line size, whereas Blocks are much larger, up to more than 100^3 nodes, to minimize communication overhead), additional data (Blocks contain management information), data layout (because of small tile size a special memory layout minimizing the number of uncoalesced transactions is needed), and hardware mapping granularity (Blocks are mapped per single process, whereas tiles are processed by separate GPU thread blocks).

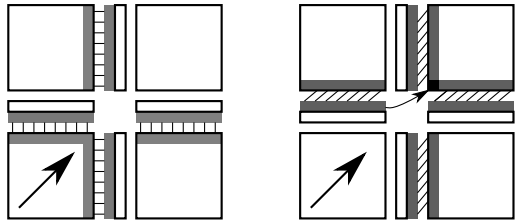

(a) write (b) read

Fig. 3. Two-step propagation toward the north-east (see arrow) of values for edge nodes in tiles with ghost buffers. Nodes and ghost buffers used in operations are marked as gray. Notice that edge values are shifted during data reads (writes are done without shift). In this way we avoid uncoalesced writes for values at the corners; only reads may be uncoalesced (black node in the right image).

3 TILES

In our implementation the whole geometry is covered by the uniform mesh of fixed-size tiles. Let a denote the number of nodes per tile edge. The number of nodes per tile is $n_{tn} = a^2$ for 2D, and $n_{tn} = a^3$ for 3D lattice arrangements. If a geometry size is not divisible by a, then the geometry is extended with solid nodes.

The tiling is implemented by the host code and done once at the geometry load. We used a very simple algorithm: first, the geometry is covered by the uniform mesh of tiles starting at node (0,0), and next the tiles containing only solid nodes are removed.

During a single LBM time step the tiles can be processed independently and in any order provided that values at the tile edges are correctly propagated. We implemented two different methods for data at tile edges synchronization.

The first method (T2C) uses two copies of f_i data and the *gather* pattern, e.g. values from neighbor nodes are accessed during the read stage. The information about tile placement is stored in an additional *tileMap* array with pointers to data for non-empty tiles (for empty tiles -1 is used). The pointers are used to find neighbor tiles needed during propagation stage. The *tileMap* array is similar to FIA but uses the tile granularity what decreases overheads proportionally to the number of nodes per tile.

In the second method (TGB) we used only one copy of f_i data and additional *ghost buffers* at tile edges. The ghost buffers contain copies of propagated f_i functions. During the propagation in a single LBM time step the values from ghost buffers can be read and written at the same time (and in any order). To prevent data races from occurring, we used two copies of ghost buffers: one copy is used for the data store, the other copy is used for the data read. After each LBM time iteration the copies are exchanged (by pointers exchange). The idea of the two-step propagation is shown in Fig. 3. For data inside tile we used the *scatter* pattern - the computed values are stored to neighbor nodes. For ghost buffers the *gather* pattern is used.

3.1 Tiles overhead

Compared with the minimal requirements defined in section 2.2, the tiles bring some memory, bandwidth and computational overheads due to: solid nodes inside tiles, additional data for storing information about tiles placement, additional data for storing node type, methods of avoiding race conditions (either two copies of f_i data or ghost buffers), uncoalesced memory transactions due to data layout, computational overhead for addresses calculation, and others. The reasons for overheads fall into one of the two classes: the overheads caused by the method and the overheads resulting from implementation. In this section we focus only on the former class. For simplification, we consider only these lattice arrangements that require only a single layer of ghost buffers (up to D2Q9 for 2D and D3Q27 for 3D geometries).

Let n_{tsn} denote the average number of solid nodes per tile, and $n_{tfn} = n_{tn} - n_{tsn}$ denote the average number of non-solid nodes per tile. We can define the average tile solidity

$$\eta_t = \frac{n_{tfn}}{n_{tn}} = \frac{n_{tn} - n_{tsn}}{n_{tn}} = 1 - \frac{n_{tsn}}{n_{tn}} = 1 - \phi_t, \qquad (17)$$

where ϕ_t is the average tile porosity, $\eta_t, \phi_t \in [0, 1]$ and $\eta_t + \phi_t = 1$. The minimum memory required to store all non-solid nodes from tile is then

$$M_{tile} = n_{tfn} \cdot M_{node} = n_{tn} \cdot \eta_t \cdot M_{node} \ \ [B] \qquad (18)$$

and the minimum transfer for tile is

$$B_{tile} = n_{tfn} \cdot B_{node} = n_{tn} \cdot \eta_t \cdot B_{node} \ \ [B]. \qquad (19)$$

3.1.1 Memory overhead

The memory overhead for tiles is a sum of four components: the overhead resulting from solid nodes inside a tile $\Delta^M_{\eta_t}$, the memory required to store node type Δ^M_{nt}, the memory used to avoid race conditions Δ^M_{rc}, and the memory needed for storing the additional tile data Δ^M_{ad}.

For all nodes in a tile (independently of the node type) all f_i values must be stored. Thus the memory overhead resulting from solid nodes in a tile is equal to

$$\Delta^M_{\eta_t} = \frac{n_{tsn} \cdot M_{node}}{M_{tile}} = \frac{n_{tsn}}{n_{tfn}} = \frac{1}{\eta_t} - 1. \qquad (20)$$

Let s_t denote a number of bytes used to store the node type (we assume that at least one byte is necessary, despite the fact that for simple cases a few bits may be enough). Since the node type is stored for all nodes inside a tile, the overhead caused by the memory required to store the node type is

$$\Delta^M_{nt} = \frac{n_{tn} \cdot s_t}{M_{tile}} = \frac{s_t}{M_{node}} \cdot \frac{1}{\eta_t}. \qquad (21)$$

The values of Δ^M_{rc} and Δ^M_{ad} depend on the method of race condition prevention.

3.1.1.1 Two copies: In this method an additional copy of all f_i values is needed for each node in a tile. The memory overhead is then

$$\Delta^M_{rc} = \frac{n_{tn} \cdot M_{node}}{M_{tile}} = \frac{1}{\eta_t}. \qquad (22)$$

The additional tile datum is only the *tileMap* array where one index per each tile (empty and non-empty) is stored. In practice, the index is $s_{ti} = 4$ byte wide what allows to use 2^{32} non-empty tiles, which require much more memory than is available on current GPUs. Let N_{tiles} denote the number

of all tiles and N_{ftiles} denote the number of non-empty tiles. The overhead resulting from memory needed to store *tileMap* array is then

$$\Delta^M_{ad} = \frac{N_{tiles} \cdot s_{ti}}{N_{ftiles} \cdot M_{tile}} = \frac{N_{tiles} \cdot s_{ti}}{N_{ftiles} \cdot \eta_t \cdot n_{tn} \cdot M_{node}}. \quad (23)$$

Notice that since $s_{ti}/(n_{tn} \cdot M_{node}) < 0.01$ even for D2Q9 lattice, single precision data and small tiles containing $n_{tn} = 4^2$ nodes, thus Δ^M_{ad} is negligible unless the ratio N_{tiles}/N_{ftiles} is large. For geometries used in this work $N_{tiles}/N_{ftiles} \in (2.3, 8.6)$.

The complete memory overhead for two copies of f_i data is then

$$\Delta^M_{T2C} = \frac{1}{\eta_t} \cdot \left(2 - \eta_t + \frac{1}{M_{node}} \cdot \left(s_t + \frac{N_{tiles}}{N_{ftiles}} \cdot \frac{s_{ti}}{n_{tn}}\right)\right). \quad (24)$$

Usually the components Δ^M_{nt} and Δ^M_{ad} can be skipped ($1/M_{node} < 0.03$ even for D2Q9 lattice and single precision data), what allows to approximate the memory overhead as $\Delta^M_{T2C} \approx (2 - \eta_t)/\eta_t$.

Equation (24) contains many parameters, but for concrete implementation most of them are constant. For example, for double precision data ($s_d = 8$), two bytes per node type ($s_t = 2$), and $s_{ti} = 4$ the memory overhead is

$$\Delta^M_{T2C} = \frac{2.028 + 0.00022 \cdot \frac{N_{tiles}}{N_{ftiles}}}{\eta_t} - 1 \quad (25)$$

for D2Q9 lattice and tiles containing 16^2 nodes. For D3Q19 lattice and 4^3 nodes per tile

$$\Delta^M_{T2C} = \frac{2.013 + 0.00041 \cdot \frac{N_{tiles}}{N_{ftiles}}}{\eta_t} - 1. \quad (26)$$

3.1.1.2 Ghost buffers: The memory needed for ghost buffers at tile edges depends on two factors. The first is the number of f_i functions propagated in directions toward the faces, the edges, and the corners. The second is the number of common faces, edges, and corners between neighbor tiles. For example, even if the current tile has only neighbors contacting either vertexes (for 2D/3D) or edges (for 3D), the additional buffers on the edges/faces must be allocated (see Fig. 3b the black node in the corner). Although for tiles contacting vertexes only the single-element ghost buffers could be allocated, to simplify the implementation we assumed that ghost buffers are allocated only in full size (n_{tn}/a values per buffer what gives a for 2D and a^2 for 3D geometries). The performance penalty caused by this simplification is small - for test cases from Table 1 in section 4 only less than 2% of allocated ghost buffers could have the reduced size.

In practice some ghost buffers are unnecessary. Even for geometries completely filled with tiles, the ghost buffers at geometry edges are not used. If the geometry is sparse and some tiles are removed, then the corresponding ghost buffers can be removed as well (Fig. 4). To take into account these removed ghost buffers, we use an additional, geometry dependent coefficient α_M that is the ratio of allocated ghost buffers to all possible ghost buffers. In general, we define only that $\alpha_M < 1$ (α_M is equal to 1 for infinite geometries without edges). For our test cases $\alpha_M > 0.7$ was

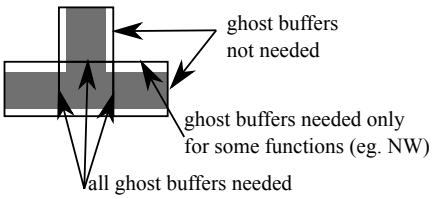

Fig. 4. If some tiles are removed, then the ghost buffers are needed only for edges between nonempty tiles. Non-solid nodes marked as gray.

computed for each case separately after the geometry tiling (see Table 1).

To prevent data races from occurring, each PDF f_i needs a set of two ghost buffers: one for a read and one for a write. Let q_s denote the number of f_i functions for which only one set of ghost buffers is needed (the functions propagated toward the edges in 2D and toward the faces in 3D geometries), q_d denote the number of f_i functions for which two sets of ghost buffers are needed (see Fig. 3, the functions propagated toward the corners in 2D and toward the edges in 3D geometries), and q_t denote the number of functions requiring three sets of ghost buffers (the functions propagated toward the corners for 3D geometries). For D2Q9 lattice we get $q_s = 4, q_d = 4, q_t = 0$, for D3Q19 lattice the values are $q_s = 6, q_d = 12, q_t = 0$, and for D3Q27 $q_s = 6, q_d = 12, q_t = 8$. The number q of all f_i functions is then $q = q_s + q_d + q_t + 1$ (the "+1" is for the not propagated function f_0). The maximum number of ghost buffers per tile is $2 \cdot (q_s + 2 \cdot q_d + 3 \cdot q_t)$.

Having these relationships, the average (per tile) real memory usage for ghost buffers is

$$M_{gb} = 2 \cdot (q_s + 2 \cdot q_d + 3 \cdot q_t) \cdot \frac{n_{tn}}{a} \cdot s_d \cdot \alpha_M \quad [B] \quad (27)$$

where s_d is defined in section 2.2 as the number of bytes for storing a single f_i value. Using M_{tile} defined in Eqn. (18) the memory overhead resulting from ghost buffers is

$$\Delta^M_{rc} = \frac{M_{gb}}{M_{tile}} = C_{gb} \cdot \frac{2 \cdot \alpha_M}{a \cdot \eta_t} \quad (28)$$

where $C_{gb} = (q_s + 2 \cdot q_d + 3 \cdot q_t)/q$ is a constant defined for given lattice arrangement. For D2Q9, D3Q19 and D3Q27 lattice arrangements the value of C_{gb} is $\frac{4}{3}$, $\frac{30}{19}$, and 2.

The last element affecting the memory overhead is the memory needed for storing the additional tile data. The majority of this memory are pointers to ghost buffers. Since pointers can be replaced with array indexes, in typical cases each pointer (index) requires $s_{gbi} = 4$ bytes that allow addressing more ghost buffers than it can be stored in current generation GPUs. Due to the *gather* pattern (data propagation during the read stage), for each tile we need to use $q_s + 2 \cdot q_d + 3 \cdot q_t$ pointers to ghost buffers that are written, and $q_s + 3 \cdot q_d + 7 \cdot q_t$ pointers to ghost buffers that are read. The overhead resulting from storing ghost buffer indexes is then

$$\Delta^M_{ad} = C_{gbi} \cdot \frac{s_{gbi}}{M_{tile}} \quad (29)$$

where $C_{gbi} = 2 \cdot q_s + 5 \cdot q_d + 10 \cdot q_t$ is a constant defined for given lattice arrangement. For the D2Q9, D3Q19 and D3Q27 lattice arrangements C_{gbi} is equal to 28, 72, and 152 indexes per tile.

The complete memory overhead for tiles with ghost buffers is then

$$\Delta^M_{TGB} = \frac{1}{\eta_t}\cdot\left(1-\eta_t+\frac{1}{M_{node}}\cdot\left(s_t+\frac{C_{gbi}\cdot s_{gbi}}{n_{tn}}\right)+ + \frac{2\cdot\alpha_M\cdot C_{gb}}{a}\right). \quad (30)$$

For double precision data ($s_d = 8$), two bytes per node type ($s_t = 2$), and $s_{gbi} = 4$ the memory overhead is

$$\Delta^M_{TGB} = \frac{1.034 + 0.167\cdot\alpha_M}{\eta_t} - 1 \quad (31)$$

for D2Q9 lattice and tiles containing 16^2 nodes. For D3Q19 lattice and 4^3 nodes per tile

$$\Delta^M_{TGB} = \frac{1.043 + 0.789\cdot\alpha_M}{\eta_t} - 1. \quad (32)$$

Notice that for the D3Q19 lattice arrangement the memory overhead Δ^M_{rc} caused by ghost buffers is almost 5 times larger than for D2Q9, mainly due to 4 times smaller tile edge.

3.1.2 Bandwidth overhead

In bandwidth overhead estimation we assume that for each node in each non-empty tile only the node type field must be read (s_t bytes per node). When a node is solid, no other operations are done, thus no unnecessary transfers of f_i values occur. Also, the f_i values for non-solid nodes are transferred from/to either tile or ghost buffers, thus for each node only a minimal amount of f_i data (see Eqn. (10)) is read/written. This way the bandwidth overhead results only from transfers of the node type field Δ^B_{nt} and transfers of additional tile data Δ^B_{ad}.

For simplification we also ignore the fact that some writes of f_i values can be skipped, when the neighbor node is solid. Additionally, since we use the same measure for all non-solid nodes (fluid and boundary), then for the whole geometry each node requires the same amount of transferred data.

The node type field must be read not only for the current node, but also for all $(q-1)$ neighbor nodes of every non-solid node to avoid propagation to/from solid nodes. However, we assume that the node type fields for a single tile are internally buffered in some way (registers/cache/scratchpad), thus for each tile the node type fields are read only once for nodes inside a tile and once for nodes forming a "halo" (1 node wide) around the tile. The value of Δ^B_{nt} is then

$$\Delta^B_{nt} = \frac{(a+2)^d\cdot s_t}{B_{tile}} \quad (33)$$

where $d \in \{2,3\}$ is the space dimension and a is the number of nodes per tile edge. The value of Δ^B_{ad} depends on the method of race conditions prevention.

3.1.2.1 Two copies: In this method each non-empty tile needs to load pointers to the neighbor tiles required for propagation, thus $q-1$ indexes of s_{ti} bytes are loaded for each non-empty tile. We do not need to load the index of the current tile because it can be computed in another way. The bandwidth overhead caused by additional data is then

$$\Delta^B_{ad} = \frac{(q-1)\cdot s_{ti}}{B_{tile}} \quad (34)$$

and the complete bandwidth overhead for T2C is

$$\Delta^B_{T2C} = \frac{1}{B_{tile}}\cdot\left((a+2)^d\cdot s_t + (q-1)\cdot s_{ti}\right). \quad (35)$$

For double precision data, two bytes per node type and four bytes per tile index ($s_{ti} = 4$) the bandwidth overhead $\Delta^B_{T2C} = \frac{0.0184}{\eta_t}$ for D2Q9 lattice with 16^2 nodes per tile and $\Delta^B_{T2C} = \frac{0.0259}{\eta_t}$ for D3Q19 lattice and 4^3 nodes per tile.

3.1.2.2 Ghost buffers: For the implementation based on ghost buffers each non-empty tile needs to load C_{gbi} indexes of s_{gbi} bytes each. Though some indexes could be skipped when all nodes placed on tile face/edge are solid, the detection of such situation seems a complex operation. Moreover, the number of skipped indexes depends on geometry; in fact, this number may be proportional to $(1-\alpha_M)$, which for cases from Table 1 is smaller than 0.3. Also, Fig. 10 shows that the value of Δ^B_{ad} was significantly lower than Δ^B_{nt} for all test geometries, thus the decrease of Δ^B_{ad} even by 30% results in much lower decrease of total overhead. Therefore, in the below considerations we can safely assume that each non-empty tile always loads all indexes to ghost buffers. The bandwidth overhead caused by additional data is then

$$\Delta^B_{ad} = \frac{C_{gbi}\cdot s_{gbi}}{B_{tile}} \quad (36)$$

and the complete overhead for TGB is

$$\Delta^B_{TGB} = \frac{1}{B_{tile}}\cdot\left((a+2)^d\cdot s_t + C_{gbi}\cdot s_{gbi}\right). \quad (37)$$

For double precision data, two bytes per node type and four bytes per ghost buffer index ($s_{gbi} = 4$) the bandwidth overhead $\Delta^B_{T2C} = \frac{0.0206}{\eta_t}$ for D2Q9 lattice with 16^2 nodes per tile and $\Delta^B_{T2C} = \frac{0.0370}{\eta_t}$ for D3Q19 lattice and 4^3 nodes per tile.

3.1.2.3 Burst transactions impact: The overhead estimates presented above define the minimal possible overheads obtainable on the ideal machine (section 2.2). For the real parallel machines the additional overheads appear.

Modern machines with DRAM memory use *burst* transactions including an aligned block of at least s_b bytes (usually $s_b \in \{32, 64, 128\}$), therefore always several neighbor values are transferred during a single memory transaction. To achieve the minimal bandwidth overhead the node data must be placed in memory in a way that allows to fully utilize all memory transactions. Since this may be difficult, we also consider a bandwidth overhead for implementation with suboptimal memory layout.

In the extreme case, each read/write of a value can induce a separate s_b–byte wide transaction. However, due to the spatial and temporal locality of data references, it seems reasonable to assume that the maximum transfer per tile contains, in addition to the node type fields and the additional tile data, all f_i values from tile and all f_i values from ghost buffers (for tiles with ghost buffers). We should also take into account that for ghost buffers containing the values from tile corners only a single transaction is needed to read a single f_i value (writes of single corner values does not occur). Notice that for 3D tiles we assume that even for ghost buffers containing only a single edge of values (a

values) the full ghost buffer with a^2 values is transferred because it does not require special data arrangement.

The bandwidth overhead resulting from transfers of full tile data is

$$\Delta^B_{ftd} = \frac{n_{tn} \cdot B_{node} - B_{tile}}{B_{tile}} = \frac{1}{\eta_t} - 1. \quad (38)$$

Let q_c denote the number of corner values that are read in separate, n_b byte transactions ($q_c = q_d$ for 2D and $q_c = q_t$ for 3D lattice arrangements). The additional bandwidth resulting from transfers of all ghost buffers contains two components: the transfers of all ghost buffers except the ones containing only single, corner value

$$B_{gbnc} = (C_{gbi} - q_c) \cdot \frac{n_{tn}}{a} \cdot s_d, \quad (39)$$

and the reads of all corner values

$$B_{gbc} = q_c \cdot n_b. \quad (40)$$

In real geometries some ghost buffers are not needed, thus the maximum transfer should be scaled by a coefficient $\alpha_B \leq 1$ (similar to α_M) defined as a ratio of actually transferred ghost buffer values to the maximum number of transferred ghost buffer values. Notice that α_B slightly differs from α_M because for tile the read ghost buffers differ from the written ghost buffers. In our sparse 2D geometries α_B was slightly lower than α_M ($\alpha_B > 0.9 \cdot \alpha_M$).

The bandwidth overhead including burst transactions impact is then

$$\Delta^B_{T2Cbt} = \Delta^B_{T2C} + \Delta^B_{ftd} \quad (41)$$

for tiles with two copies and

$$\Delta^B_{TGBbt} = \Delta^B_{TGB} + \Delta^B_{ftd} + \frac{B_{gbnc} + B_{gbc}}{B_{tile}} \cdot \alpha_B \quad (42)$$

for tiles with ghost buffers. Both values defined in Eqn. (41) and (42) may be treated as the pessimistic estimates of bandwidth overhead.

3.2 Implementation details

The code was written in CUDA C language ver. 7.5, the first version with official support for many C++11 features. The support of C++11 allowed us to design a heavily templated generic kernel code that can be specialized for fluid and collision models, tile size, and enabled optimizations. The main disadvantage of this solution is a long compilation time.

In our implementation a single GPU thread block is responsible for processing of a single tile what allows for easy synchronization of computations within the tile. Consecutive GPU threads are assigned to nodes from tile using row order. To minimize the number of memory transfers, we combined collision, propagation (streaming) and boundary computations for a single node into a single GPU kernel. The node data (f_i, v, ρ, node type) are transferred only once and local copies are stored in registers and in the shared memory.

Figures 5 and 6 show the two versions of the kernel. We started from the version of TGB for D2Q9 lattice arrangement and 16^2 nodes per tile. Then we prepared the version that uses T2C for D3Q19 lattice arrangement and 4^3 nodes per tile, where we also fixed issues observed in the D2Q9 version. The main difference between TGB and T2C implementations is in the code responsible for propagation step.

```
1: load node type
2: if node not solid then
3:    load all f_i
4: end if
5: load ghost buffers {2 BARRIERS due to WLP}
6: if node not solid then
7:    process boundary nodes {calculate u, v, ρ}
8:    if node is boundary then
9:       store all f_i
10:   end if
11:   collide
12: end if
13: store ghost buffers {1 BARRIER due to WLP}
14: if node not solid then
15:   scatter all f_i
16: end if
```

Fig. 5. Structure of the GPU kernel implementing single LBM time iteration for a single node within the tile with ghost buffers containing 16^2 nodes. WLP denotes warp level programming.

```
1: load copy of tile bitmap {3^3 values}
2: load node types from current tile {4^3 values}
3: load node types from neighbor tiles {WLP}
4: BARRIER
5: if node not solid then
6:    load f_0
7:    for i ∈< 1..18 > do
8:       compute address of neighbor node in direction i
9:       if neighbor node not solid then
10:         gather f_i from neighbor node
11:      end if
12:   end for{all f_i copied to registers}
13:   process boundary {calculate v, ρ}
14:   collide
15:   store all f_i {all coalesced}
16: end if
```

Fig. 6. Structure of the GPU kernel implementing single LBM time iteration for a single node within the tile containing 4^3 nodes with two copies of f_i data.

The propagation between two nodes is done only when the target and source nodes are not solid. In both implementations we could use the node type fields for nodes inside the tile, but nodes from neighbor tiles had to be treated in a different way. For TGB a partial information about the node type for nodes from neighbor tiles is stored in the values put into the ghost buffers - if the node corresponding to the given value in the ghost buffer is solid, then the value in the ghost buffer is set as the predefined not a number (NaN). In T2C version the node type fields from neighbor tiles are tested directly. In both versions we used shared memory to store either ghost buffers propagated horizontally (line 5 in Fig. 5) or node types from current and neighbor tiles (lines 2–3 in Fig. 6). This allows minimizing the memory bandwidth usage caused by multiple reads of the same

values and uncoalesced memory accesses. However, since the number of either ghost buffers or node type fields for current tile with an additional halo does not map well to GPU threads, we used a warp level programming (WLP) to balance the load of warps assigned to a tile. Though WLP requires additional synchronization in the kernel, it is amortized with interest - in the kernels without WLP both the performance and the real GPU utilization were lower. The unfavorable issue of the WLP based solution is a low code portability, because each change of the tile size, thread to nodes mapping, or warp size requires a new version of the propagation code.

Besides the method of providing information about node types from neighbor tiles, the second difference between the TGB/T2C implementations of propagation is the data access pattern. In the TGB version the propagation is implemented as a scattering data to neighbor nodes and is additionally partitioned into two separated steps because the propagation of f_i values for nodes at tile edges requires a sequence of writes and the following reads to/from ghost buffers (see Fig. 3). Notice that both these propagation steps must be done before the collision, and that the synchronization between the first and the second step is required to prevent data races for nodes at tile edges (because the tiles are processed in parallel). Thus, our single iteration starts with the second step of propagation (line 5 in Fig. 5) and performs the first propagation step at the iteration end (lines 13–16 in Fig. 5).

Additionally, for the propagation done as the data scattering some f_i values can be not correctly updated from neighbor nodes (for example when these nodes are solid). To avoid this, lines 8–10 in Fig. 5 are responsible for writing to global memory the values for these f_i functions that will not be overwritten during propagation.

For the T2C kernel we implemented the propagation as a gathering data from neighbor nodes. This allowed us to simplify the code because no equivalent of lines 8–10 from Fig. 5 was needed. Due to a lack of accesses to ghost buffers, we could also remove synchronization points during writes (all writes are only to nodes within a tile and are perfectly coalesced), what allowed for much simpler, regular code with the single barrier at line 4.

The last important difference between kernels TGB and T2C is in the neighbor node address computations (line 8 in Fig. 6, not shown in Fig. 5). In the TGB version the address of neighbor node data depends only on a location of the current node and can be determined based only on current tile data. The T2C version additionally needs locations of neighbor tiles, thus the two-step procedure is needed. First, the index of the tile containing the neighbor node is computed - we used the values from $\{-1, 0, 1\}$ due to the local copy of the tile bitmap (created in line 1). Then, when the neighbor tile is not empty, a few indexes for the neighbor node are computed since we need both information about the neighbor node type (this is gathered from the copy in shared memory) and the value of proper f_i (which is stored in global memory).

We also applied a set of widely know optimizations: loop unrolling to increase instruction level parallelism due to interlacing of independent streams of instructions, using shared memory to decrease register pressure, tuning the usage of registers for each specialization of the kernels, minimizing cost of divergent branches by placing time-consuming memory accesses outside the divergent code, replacing divisions with multiplications by inverse, avoiding transfers of v, ρ values, and the use of inline methods and compiler pragmas to connect clean, structured code with high performance. Additionally, for the T2C kernel we used an optimized data layout in memory for each array of f_i values to minimize the number of uncoalesced memory transactions. After this optimization, the real number of memory transactions for each completely filled tile was only about 1% higher than the minimum value defined by Eqn. (19) due to a minimized number of uncoalesced reads and an increased cache utilization. The address calculations are done by *constexpr* methods.

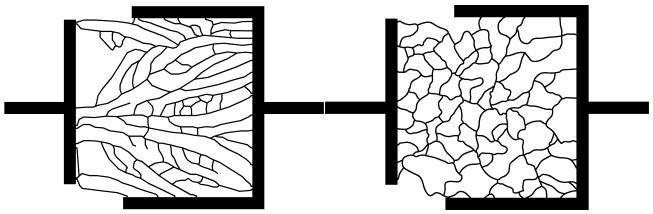

(a) ChipA (b) ChipB

Fig. 7. Geometries for 2D test cases. Non-solid nodes are marked as black. Channel width is defined for narrow channels in the middle of the drawing.

4 Results

All simulations were run in double precision on a computer with the NVIDIA GTX TITAN device (Kepler architecture) clocked at 823 MHz with 6 GB GDDR5 memory clocked at 3.004 GHz, Intel i7-3930K CPU and 64 GB 4-channel DDR3 DRAM. Simulations for D3Q19 lattice arrangement were done using the T2C kernel, for D2Q9 we used the TGB kernel.

4.1 Test geometries

The simulations were run for dense and sparse 2D and 3D geometries. The dense geometries consisted of a fluid flow in a square chamber with a moving lid. The sparse geometries for the D3Q19 lattice arrangement were prepared on the base of cases similar to presented in [16] (three arrays of randomly arranged spheres with the diameter of 40 lattice units and the blood flow in the cerebral aneurysm) and [17] (aorta with coarctation). For the D2Q9 lattice arrangement we used the microvascular structures from [20] (Fig. 7 ChipA) and [21] (ChipB).

The parameters of the cases are shown in Table 1. The tile size is 16^2 nodes for 2D and 4^3 nodes for 3D. The ratio of boundary to all non-solid nodes varies from 0.05 (for the cerebral aneurysm) up to 0.5 (for RAS_0.7). The porosity of geometries varies between 0.7 and 0.91.

4.1.1 Tile parameters

The value of the most important tile parameter, the tile solidity η_t, depends mainly on other factors than the geometry solidity η. This is especially visible for 2D cases

TABLE 1
Cases used for tests. N_{nodes} - number of all nodes, η - geometry solidity defined by Eqn. (12), η_t - average tile solidity defined by Eqn. (17), Δ^M, Δ^B - memory and bandwidth overheads for tiles with two copies (T2C), tiles with ghost buffers (TGB), CM, and FIA, and α_M - coefficient defined in section 3.1.1.2. The suffixes after *ChipA*/*ChipB* names denote channel width in nodes. The suffixes after *RAS* (random arranged spheres) names denote porosity.

Case	N_{nodes}	η	η_t	α_M	Δ^M				Δ^B			
					TGB	T2C	FIA	CM	TGB	T2C	FIA	CM
ChipA_08	2 392 350	0.21	0.58	0.80	**1.01**	2.49	1.27	1.44	0.035	**0.032**	1.133	0.22
ChipB_08	2 587 417	0.20	0.60	0.82	**0.94**	2.37	1.27	1.44	0.034	**0.031**	1.137	0.22
ChipA_16	9 574 320	0.20	0.71	0.86	**0.65**	1.85	1.27	1.44	0.029	**0.026**	1.137	0.22
ChipB_16	10 370 560	0.20	0.74	0.87	**0.58**	1.73	1.28	1.44	0.028	**0.025**	1.141	0.22
ChipA_32	38 075 466	0.20	0.83	0.91	**0.43**	1.45	1.28	1.44	0.025	**0.022**	1.139	0.22
ChipB_32	41 455 800	0.20	0.84	0.92	**0.42**	1.42	1.28	1.44	0.025	**0.022**	1.142	0.22
RAS_0.7	192^3	0.30	0.70	0.86	1.46	1.88	**1.09**	1.47	0.053	**0.037**	1.044	0.24
RAS_0.8	192^3	0.20	0.63	0.82	1.70	2.22	**1.13**	1.47	0.059	**0.041**	1.066	0.24
RAS_0.9	192^3	0.10	0.51	0.76	2.21	2.94	**1.27**	1.47	0.072	**0.051**	1.133	0.24
Aneurysm	84 607 488	0.18	0.93	0.97	**0.95**	1.17	1.15	1.47	0.040	**0.028**	1.075	0.24
Coarctation	6 990 336	0.09	0.81	0.91	**1.19**	1.50	1.28	1.47	0.046	**0.032**	1.140	0.24

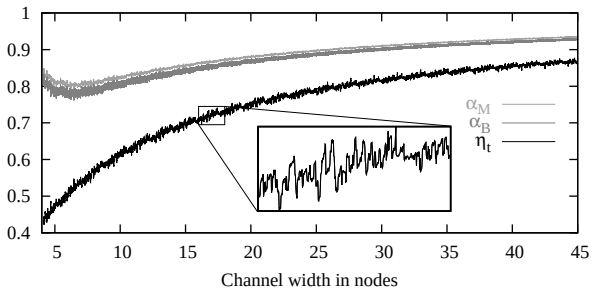

Fig. 8. Tile parameters for ChipA geometry.

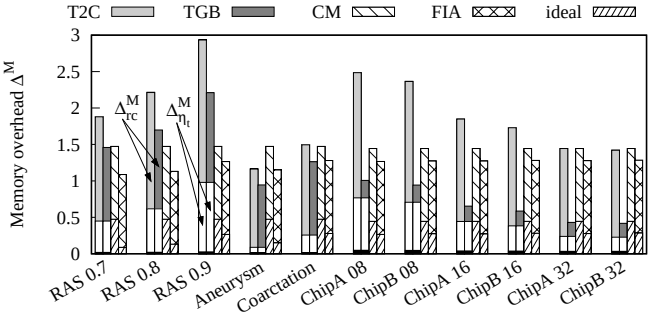

Fig. 9. Memory overhead estimates for tiles, CM and FIA implementations. Small black rectangles (almost invisible) denote Δ^M_{nt} and Δ^M_{ad}.

where different values of η_t are obtained for the same η. Also, for the aneurysm and coarctation cases η_t becomes high despite the low η. Fig. 8 shows tile parameters as a function of the number of nodes per channel width. Only ChipA is analyzed, but for ChipB the values are similar. The value of η_t greatly depends on the channel width: η_t becomes larger than 0.8 for 27 node wide channels ($1.7\times$ tile edge) and larger than 0.85 for 39 node wide channels ($2.4\times$ tile edge).

The values of α_M and α_B are between 0.76 and 0.94. Both coefficients are close to each other (the difference was below 0.073) and α_B is usually lower than α_M.

4.1.2 Memory overhead

The comparison of memory overhead estimates for sparse test geometries is shown in Table 1 and Fig. 9, which also shows the components of both Δ^M_{T2C} and Δ^M_{TGB} and the values of Δ^M for "ideal" implementations of CM and FIA. The results show that for 2D geometries with a high tile solidity the tiles with ghost buffers allow for significant reduction of memory usage. In other cases the memory reduction is more difficult to achieve.

Though for D3Q19 lattice arrangement T2C has the memory overhead similar (about 20–30% higher) to TGB, for D2Q9 the memory overhead for T2C is 2.5–3.4 times higher than for TGB. This results from a low overhead for TGB due to large tiles and a small number of ghost buffers needed for D2Q9, and allows achieving the lowest memory overhead for all tested 2D geometries. For both ChipA_32 and ChipB_32 geometries the total memory usage for FIA/CM methods is 1.6–1.7 times higher than for TGB ($2.28/1.43$ and $2.44/1.43$). Less favorable results are for D3Q19 lattice arrangement - the memory overhead for tiles is either close to or higher than for FIA/CM. TGB has the lowest memory overhead only for the geometries with the high tile solidity (aneurysm and coarctation), and the gains are much lower than for D2Q9. Memory overhead for T2C is always higher than for FIA. Compared with CM, T2C requires less additional memory only for the aneurysm case.

For all 3D cases, and for T2C on D2Q9 lattice, the majority of the tile memory overhead is Δ^M_{rc} (either the second copy of f_i or ghost buffers). Only for TGB on D2Q9 lattice the memory overhead is close to $\Delta^M_{\eta_t}$, which can be treated as an overhead for "ideal" implementation of tiles. Removal of Δ^M_{rc} could significantly decrease memory requirements for tiles to values lower than for implementations of CM/FIA. However, if we compare $\Delta^M_{\eta_t}$ with overheads for "ideal" implementations of CM/FIA, the tiles have lower memory usage only for cases with high tile solidity, and the gains are small.

4.1.3 Bandwidth overhead

The comparison of memory bandwidth overhead estimates for sparse test geometries is shown in Table 1 and Fig. 10. For all tested geometries the overhead estimates for both versions of tile processing are significantly lower than for

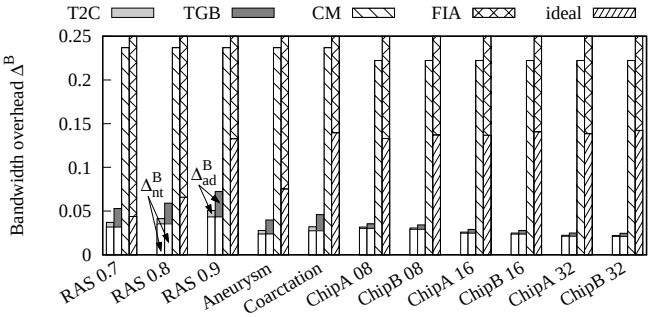

Fig. 10. Bandwidth overhead estimates for sparsity handling techniques. The overheads for FIA exceed 1.0 for all cases.

CM/FIA, up to $10\times$ than CM and up to $50\times$ than FIA. Also, the overhead estimates for tiles are lower than the ideal estimates for CM/FIA for all cases but RAS_0.7. The bandwidth overhead estimates for the ideal implementation for tiles (equal to Δ^B_{nt}) are always lower than the ideal estimates for CM/FIA. These results show that the tiles allow achieving the highest performance for sparse geometries. The good implementations for tiles could allow achieving the performance for presented cases only a few percent lower than the peak performance for dense geometries.

The bandwidth overhead for tiles is determined by two components: Δ^B_{nt} and Δ^B_{ad}. In all cases Δ^B_{ad} is smaller than Δ^B_{nt}, thus performance optimizations should focus on minimizing transfers of the node type fields. However, since for all test cases the overall bandwidth overhead is lower than 0.1 (usually close to 0.05), it can be difficult to decrease it further.

4.2 Performance comparison

To provide some measure allowing for comparison of different implementations for different lattice arrangements and on different machines, we calculated the number of bytes transferred per node from Eqn. (10) and the minimum required memory bandwidth to achieve the reported performance P [MLUPS] as $B_{case} = P \cdot B_{node}$. For sparse geometries we took into account the transfers only for non-solid nodes. Next, we compared this value with the maximum theoretical GPU bandwidth B_{GPU} and calculated a bandwidth utilization

$$BU = \frac{B_{case}}{B_{GPU}} \tag{43}$$

shown in Table 2. This measure is based on [18] and is similar to measures used in many papers (for example [6], [7], [11]). Due to its strict assumptions, it shows how the given implementation is close to the ideal one for bandwidth bound codes. According to estimations shown in section 2.2, the LBM implementations on GPU are bandwidth bound except either the implementations of complex models on the previous generations GPUs or the lattice arrangements with a high number of links [22]. Notice only that we used the theoretical maximum of memory bandwidth instead of the measured.

4.2.1 Dense geometries

Performance comparison of different LBM implementations for dense geometries is shown in Table 2. We omitted results when performance was limited by the low computational power of the GPU, for example double precision computations on GTX280 [5], which resulted in $BU < 0.3$.

For D3Q19 lattice arrangement and double precision, only [10] had $1.17\times$ higher BU than the tile-based solution. Our implementation had BU better than [6] ($1.17\times$) and than [14] ($1.4\times$). Also many implementations for single precision had lower BU than our version. Since we compared our solution designed for sparse geometries with highly optimized implementations for dense geometries (without overheads), these results show a low additional cost imposed by tiles. Notice that performance of our implementation is decreased not only by bandwidth overhead defined in Eqn. (35) for $\eta_t = 1$, but also by uncoalesced memory transactions resulting from data placement inside tiles.

The results for D2Q9 lattice arrangement were less favorable for our implementation. Despite the higher ratio of bandwidth to computations (section 2.2) and the lower register pressure due to about two times smaller size of data per node, the tile based implementation for D2Q9 had up to $1.36\times$ lower BU than for D3Q19 lattice. This low performance results from many synchronization points and a high complexity of code for ghost buffers handling. After removing the barriers the performance for D2Q9 lattice and BGK incompressible model increased to 1300 MLUPS ($BU = 0.649$). When we completely removed the code responsible for ghost buffers handling, the performance was 1615 MLUPS ($BU = 0.806$). The second reason of a low performance is a lack of optimized memory layout in kernels for the D2Q9 lattice.

Table 2 shows also an additional, interesting phenomena: though in theory the floating point calculations should be completely masked by memory transfers [4], the performance of our code decreased for more complex collision and fluid models. This results from explicit barriers in our code - after removing all synchronization points from the D2Q9 BGK-incompressible kernel the performance was the same for versions with and without collision computations. The decrease of performance for more complex models is also caused by a lower GPU occupancy resulting from increased register usage (especially visible for MRT on D3Q19 lattice).

4.2.2 Sparse geometries

Performance comparison for sparse 3D geometries is shown in Table 3. For the CM method we placed only results from [16], which are for optimized implementation of [15] on more recent hardware.

The results confirm theoretical analysis showing that the performance of our solution strongly depends on the tile solidity. For the geometry with low tile solidity ($\eta_t = 0.51$ for RAS_0.9) our implementation achieved BU 13% lower than [16]. However, when the tile solidity increased, the BU value for the tile-based solution became significantly higher than for [16] (up to $1.49\times$ for the aneurysm case).

Although for the aneurysm case we compared our single-GPU implementation with the multi-GPU version from [16], which requires additional communication, the

TABLE 2
Performance comparison for dense geometries. For [3] the collision/fluid model is deduced from information in the paper.

work	model	MLUPS	BU	lattice	precision	GPU
[10]	BGK q-compr.	649	0.790	D3Q19	DP	Tesla K20x
[6]	BGK incompr.	292	0.616	D3Q19	DP	Tesla C2070
[14]	BGK q-compr.	234	0.480	D3Q19	DP	Tesla M2050
[1]	BGK q-compr.	1036	0.757	D3Q19	SP	Tesla K20c
[4]	BGK q-compr.	512	0.696	D3Q19	SP	GTX 295
[4]	MRT incompr.	516	0.701	D3Q19	SP	GTX 295
[7]	BGK q-compr.	400	0.543	D3Q19	SP	GTX 260
[3]	BGK q-compr.	300	0.528	D3Q19	SP	8800 GTX
[9]	MRT incompr.	443	0.455	D3Q19	SP	Tesla M2070
[8]	BGK q-compr.	375	0.321	D3Q19	SP	GTX 480
[11]	BGK q-compr.	874	0.656	D2Q9	SP	K5000m
[5]	BGK q-compr.	947	0.481	D2Q9	SP	GTX 280
[2]	BGK incompr.	670	0.465	D2Q9	SP	8800 Ultra
this	BGK incompr.	682	0.719	D3Q19	DP	GTX Titan
this	BGK q-compr.	639	0.674	D3Q19	DP	GTX Titan
this	MRT incompr.	473	0.499	D3Q19	DP	GTX Titan
this	MRT q-compr.	476	0.502	D3Q19	DP	GTX Titan
this	BGK incompr.	1060	0.529	D2Q9	DP	GTX Titan
this	BGK q-compr.	1020	0.509	D2Q9	DP	GTX Titan
this	MRT incompr.	920	0.459	D2Q9	DP	GTX Titan
this	MRT q-compr.	865	0.432	D2Q9	DP	GTX Titan

TABLE 3
Performance comparison for sparse 3D geometries. The values in the last row are estimated based on results presented in [17].

Case	Model	this work		other work						
		MLUPS	BU	work	method	MLUPS	BU	lattice	precision	GPU
RAS_0.7	BGK q-compr.	518	**0.546**	[16]	CM	334	0.488	D3Q19	DP	Tesla K20
RAS_0.8	BGK q-compr.	477	**0.503**	[16]	CM	330	0.482	D3Q19	DP	Tesla K20
RAS_0.9	BGK q-compr.	407	0.429	[16]	CM	337	**0.493**	D3Q19	DP	Tesla K20
Aneurysm	BGK q-compr.	572	**0.603**	[16]	CM	1090	0.404	D3Q19	SP	4 $\times$ Tesla C1060
Coarctation	BGK q-compr.	574	**0.605**	[17]	FIA	~ 150	~ 0.2	D3Q15	DP	GTX 680

communication overhead in [16] is low (for single precision and large geometries BU for $4\times$ Tesla C1060 version is usually ~ 0.9 of single GPU BU, e.g. 1016 vs 288 MLUPS). Even if we assume that the single-GPU version of [16] achieves $1/0.8 = 1.25\times$ higher BU, the value of BU for our implementation is still $1.19\times$ higher. Moreover, for the aneurysm we compared our double precision implementation (more difficult to achieve high performance) with the single precision version from [16].

The values from tables 2 and 3 show that the performance penalty for sparse geometries is much higher than it is apparent from Eqn. (35). It results, inter alia, from an additional memory traffic due to burst transactions partially filled with data for non-solid nodes. In practice, the bandwidth overhead for tiles with high porosity is between values defined in Eqn. (35) and (41). It is especially visible for the RAS_0.9 case where despite the low bandwidth overhead shown in Fig. 10 the real performance for T2C is low, but much higher than it results from $\Delta^B_{T2Cbt} = 1$. These observations show that the performance of our code could be improved by a better memory layout tuned for porous tiles.

5 Conclusions

In this paper, the GPU based LBM implementation for fluid flows in sparse geometries was presented. In contrary to the previous propositions for sparse geometries, which have been using the indirect node addressing, our solution covers the geometry with the uniform mesh of square tiles. Two implementations were shown: tiles with ghost buffers (TGB) and a single copy of probability distribution functions (PDF), and tiles with two copies of PDF values (T2C). For T2C we achieved up to 682 MLUPS (71.9% utilization of maximum theoretical memory bandwidth) on GTX Titan for D3Q19 lattice, double precision data, and BGK incompressible model.

For the presented method we provided the detailed theoretical model that allows to analyze the memory and bandwidth overheads for different tile sizes, machine constraints (e.g. memory transaction size) and geometry layouts (defined for example by factors describing the solidity of both tiles and geometry). The model was used to compare the overheads introduced by tiles with overheads of other sparse LBM implementations. It can also be used to determine the "quality" of implementation and to find code inefficiencies through the comparison of real performance with theoretical limits.

Based on theoretical estimates, for presented cases tiles with two copies required only 2–5% more data traffic than implementations for dense geometries. This is a significant (up to $10\times$ over CM and up to $50\times$ over FIA) reduction of an amount of ancillary data transfers required to handle geometry sparsity. Although real implementation brings additional memory transfers due to uncoalesced transactions, for D3Q19 lattice arrangement tiles with two copies allowed to achieve up to $1.49\times$ higher device memory bandwidth utilization than CM technique (for the aneurysm case) and

TABLE 4
Performance for sparse 2D geometries.

Case	BGK incompr.		BGK q-compr.		MRT incompr.		MRT q-compr.	
	MULPS	BU	MLUPS	BU	MLUPS	BU	MLUPS	BU
ChipA_08	673	0.336	659	0.329	564	0.282	578	0.289
ChipB_08	697	0.348	682	0.341	587	0.293	600	0.300
ChipA_16	842	0.420	822	0.414	705	0.352	722	0.360
ChipB_16	861	0.430	842	0.420	733	0.366	751	0.375
ChipA_32	956	0.477	943	0.471	832	0.415	849	0.424
ChipB_32	961	0.480	938	0.468	832	0.415	850	0.424

up to 3× higher than FIA (for the coarctation case). The performance of our implementation was close even to the fastest implementations for dense geometries - we achieved 85% of bandwidth utilization from [10].

Tiles can also reduce memory usage compared with FIA/CM implementations, but large gains are possible only for the lattice arrangements with a small number of links (D2Q9) and the tiles with ghost buffers. For our D2Q9 test cases the estimated memory usage for FIA/CM techniques was up to 1.6–1.7 times higher than for tiles with ghost buffers. In other cases, the use of tiles brought small reductions of memory usage only when tiles contained a small number of non-solid nodes. Moreover, in many cases tiles with 2 copies of data required more memory than CM/FIA due to solid nodes inside tiles.

For 3D geometries tiles with ghost buffers had memory overhead similar to tiles with two copies. Since ghost buffers require more complex and slower code than two copies, thus for 3D geometries the T2C version is usually a better solution.

The main drawbacks of the presented method are a rapid decrease in performance and increase of memory usage when the tiles contain solid nodes, and a complex code, especially for tiles with ghost buffers, which require many synchronization points hindering performance.

Future work includes multi-GPU version that will allow to quickly simulate the large geometries. We are also going to improve performance and memory usage by memory layouts tuned for tiles with high porosity and better tiling algorithms, which should allow to achieve a higher tile solidity.

Acknowledgments

Authors acknowledge the financial support from the National Science Centre grant no. N N501 042140 and from the Wroclaw University of Science and Technology, Faculty of Electronics, Chair of Computer Engineering statutory funds.

References

[1] M. J. Mawson and A. J. Revell, "Memory transfer optimization for a lattice Boltzmann solver on Kepler architecture nVidia GPUs," *Computer Physics Communications*, vol. 185, no. 10, pp. 2566 – 2574, 2014.

[2] J. Tölke, "Implementation of a Lattice Boltzmann kernel using the Compute Unified Device Architecture developed by nVIDIA," *Computing and Visualization in Science*, vol. 13, no. 1, p. 29, 2008.

[3] P. Bailey, J. Myre, S. Walsh, D. Lilja, and M. Saar, "Accelerating lattice Boltzmann fluid flow simulations using graphics processors," in *Parallel Processing, 2009. ICPP '09. International Conference on*, Sept 2009, pp. 550–557.

[4] C. Obrecht, F. Kuznik, B. Tourancheau, and J.-J. Roux, "A new approach to the lattice Boltzmann method for graphics processing units," *Computers & Mathematics with Applications*, vol. 61, no. 12, pp. 3628 – 3638, 2011.

[5] F. Kuznik, C. Obrecht, G. Rusaouen, and J.-J. Roux, "LBM based flow simulation using GPU computing processor," *Computers & Mathematics with Applications*, vol. 59, no. 7, pp. 2380 – 2392, 2010.

[6] J. Habich, C. Feichtinger, H. Köstler, G. Hager, and G. Wellein, "Performance engineering for the lattice Boltzmann method on GPGPUs: Architectural requirements and performance results," *Computers & Fluids*, vol. 80, no. 0, pp. 276 – 282, 2013.

[7] P. Rinaldi, E. Dari, M. Vénere, and A. Clausse, "A lattice-Boltzmann solver for 3D fluid simulation on GPU," *Simulation Modelling Practice and Theory*, vol. 25, no. 0, pp. 163 – 171, 2012.

[8] M. Astorino, J. Sagredo, and A. Quarteroni, "A modular lattice Boltzmann solver for GPU computing processors," *SeMA Journal*, vol. 59, no. 1, pp. 53–78, 2012.

[9] C. Obrecht, F. Kuznik, B. Tourancheau, and J.-J. Roux, "Scalable lattice Boltzmann solvers for CUDA GPU clusters," *Parallel Computing*, vol. 39, no. 6–7, pp. 259 – 270, 2013.

[10] M. Januszewski and M. Kostur, "Sailfish: A flexible multi-GPU implementation of the lattice Boltzmann method ," *Computer Physics Communications*, vol. 185, no. 9, pp. 2350 – 2368, 2014.

[11] M. J. Mawson, "Interactive fluid-structure interaction with many-core accelerators," Ph.D. dissertation, The University of Manchester, Manchester, UK, 2014.

[12] C. Feichtinger, J. Habich, H. Köstler, G. Hager, U. Rüde, and G. Wellein, "A flexible Patch-based lattice Boltzmann parallelization approach for heterogeneous GPU–CPU clusters," *Parallel Computing*, vol. 37, no. 9, pp. 536 – 549, 2011.

[13] C. Feichtinger, J. Habich, H. Köstler, U. Rüde, and T. Aoki, "Performance modeling and analysis of heterogeneous lattice Boltzmann simulations on CPU–GPU clusters," *Parallel Computing*, vol. 46, pp. 1 – 13, 2015.

[14] F. Schornbaum and U. Rüde, "Massively parallel algorithms for the lattice Boltzmann method on nonuniform grids," *SIAM Journal on Scientific Computing*, vol. 38, no. 2, pp. C96–C126, 2016.

[15] M. Bernaschi, M. Fatica, S. Melchionna, S. Succi, and E. Kaxiras, "A flexible high-performance lattice Boltzmann GPU code for the simulations of fluid flows in complex geometries," *Concurr. Comput. : Pract. Exper.*, vol. 22, no. 1, pp. 1–14, Jan. 2010.

[16] C. Huang, B. Shi, Z. Guo, and Z. Chai, "Multi-GPU based lattice Boltzmann method for hemodynamic simulation in patient-specific cerebral aneurysm," *Communications in Computational Physics*, vol. 17, no. 4, pp. 960–974, 2015.

[17] C. Nita, L. Itu, and C. Suciu, "GPU accelerated blood flow computation using the lattice Boltzmann method," in *High Performance Extreme Computing Conference (HPEC), 2013 IEEE*, Sept 2013, pp. 1–6.

[18] G. Wellein, T. Zeiser, G. Hager, and S. Donath, "On the single processor performance of simple lattice Boltzmann kernels," *Comput. Fluids*, vol. 35, no. 8-9, pp. 910–919, Sep. 2006.

[19] G. Hager and G. Wellein, *Introduction to High Performance Computing for Scientists and Engineers*, 1st ed. Boca Raton, FL, USA: CRC Press, Inc., 2010.

[20] R. D. M. Travasso, E. Corvera Poiré, M. Castro, J. C. Rodrguez-Manzaneque, and A. Hernández-Machado, "Tumor angiogenesis and vascular patterning: A mathematical model," *PLoS ONE*, vol. 6, no. 5, p. e19989, May 2011.

[21] K. Y. Lin, M. Maricevich, N. Bardeesy, R. Weissleder, and U. Mahmood, "In vivo quantitative microvasculature phenotype imaging of healthy and malignant tissues using a fiber-optic confocal laser

microprobe," *Translational Oncology*, vol. 1, no. 2, p. 84–94, Jul. 2008.

[22] E. Calore, A. Gabbana, J. Kraus, S. F. Schifano, and R. Tripiccione, "Performance and portability of accelerated lattice Boltzmann applications with OpenACC," *Concurrency and Computation: Practice and Experience*, vol. 28, no. 12, pp. 3485–3502, 2016, cpe.3862.

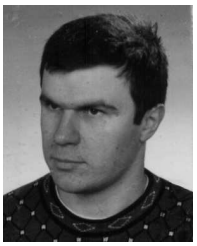

Tadeusz Tomczak received M.S. and Ph.D. degrees in computer science all from the Institute of Computers, Control, and Robotics, Wroclaw University of Technology, Poland, in 2002 and 2007 respectively. His research interests include residue number systems, fast computational hardware and parallel computing.

Roman Grzegorz Szafran received M.S. and Ph.D. degrees in chemical engineering all from the Faculty of Chemistry, Department of Chemical Engineering, Wroclaw University of Technology, Poland, in 1999 and 2004 respectively. His research interests include CFD/lattice-Boltzmann modeling of flow hydrodynamics in microfluidic devices, BIO-MEMS, multiphase flows, coating and encapsulation.